

06-19-00

A

06/16/00

Please type a plus sign (+) inside this box → ☐

PTO/SB/03 (4/98)  
 Approved for use through 08/30/2000. OMB 0851-0032  
 Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

# UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))

Attorney Docket No. CISCO-2369

First Inventor or Application Identifier Douglas Gourlay

Title Content Routing Services Protocol

Express Mail Label No. EL575422828US

U.S. PRO

09/595774

06/16/00

## APPLICATION ELEMENTS

See MPEP chapter 800 concerning every patent application contents.

## ADDRESS TO:

Assistant Commissioner for Patents  
 Box Patent Application  
 Washington, DC 20231

1. ☐ Fee Transmittal Form (e.g., PTO/SB/17)  
 (Submit an original and a duplicate for fee processing)
2. ☒ Specification (Total Pages 
 (preferred arrangement set forth below)
- Descriptive title of the invention
  - Cross References to Related Applications
  - Statement Regarding Fed sponsored R & D
  - Reference to Microfiche Appendix
  - Background of the invention
  - Brief Summary of the invention
  - Brief Description of the Drawings (if filed)
  - Detailed Description
  - Claim(s)
  - Abstract of the Disclosure
3. ☒ Drawing(s) (35 U.S.C. 113) (Total Sheets 

4. Oath or Declaration (Total Pages 
  - a. ☐ Newly executed (original or copy)
  - b. ☐ Copy from a prior application (37 C.F.R. § 1.63(d))  
 (for continuation/divisional with Box 16 completed)
  - c. ☐ DELETION OF INVENTOR(S)  
 Signed statement attached deleting  
 inventor(s) named in the prior application,  
 see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).

5. ☐ Microfiche Computer Program (Appendix)
6. Nucleotide and/or Amino Acid Sequence Submission  
 (if applicable, all necessary)
- a. ☐ Computer Readable Copy
  - b. ☐ Paper Copy (identical to computer copy)
  - c. ☐ Statement verifying identity of above copies

## ACCOMPANYING APPLICATION PARTS

7. ☐ Assignment Papers (cover sheet & document(s))
8. ☐ 37 C.F.R. § 3.73(b) Statement ☐ Power of  
 (when there is an assignee) Attorney
9. ☐ English Translation Document (if applicable)
10. ☐ Information Disclosure ☐ Copies of IDS  
 Statement (IDS/PTO-1449 Citations)
11. ☐ Preliminary Amendment
12. ☒ Return Receipt Postcard (MPEP 503)  
 (Should be specifically itemized)
13. ☐ Small Entity ☐ Statement filed in prior application,  
 Statement(s) Status still proper and desired  
 (PTO/SB/05-12)
14. ☐ Certified Copy of Priority Document(s)  
 (if foreign priority is claimed)
15. ☒ Other: Transmittal Letter

NOTE: FOR INVENTORS TO BE ENTITLED TO PAY SMALL ENTITY  
 FEES, A SMALL ENTITY STATEMENT IS REQUIRED ON C.F.R. § 1.53(b)(2) EXCEPT  
 IF ONE FILED A PRIOR APPLICATION IS RELIED UPON OF C.F.R. § 1.53(b)(2)

16. If a CONTINUING APPLICATION, check appropriate box, and supply the requisite information below and in a preliminary amendment
- ☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No. \_\_\_\_\_
- Prior application information: Examiner \_\_\_\_\_ Group / Art Unit: \_\_\_\_\_

For CONTINUATION or DIVISIONAL APPS only: The entire disclosure of the prior application, from which an oath or declaration is supplied  
 under Box 4b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by  
 reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

## 17. CORRESPONDENCE ADDRESS

☐ Customer Number or Bar Code Labelor ☐ Correspondence address below

(Insert Customer No. or Attach Bar code label here)

Name	David B. Ritchie				
	D'Alessandro & Ritchie				
Address	P.O. Box 640640				
City	San Jose	State	CA	Zip Code	95164-0640
Country	USA	Telephone	408-441-1100	Fax	408-441-8400

Name (Print/Type)	Mary S. Hanish	Registration No. (Attorney/Agent)	42,626
Signature		Date	6/16/00

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any  
 comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office,  
 Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents,  
 Box Patent Application, Washington, DC 20231.

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

In re Application of:	) Art Unit:
Douglas Gourlay	) Examiner:
Serial No. [Not yet assigned]	)
Filed: June 16, 2000	)
For: <u>CONTENT ROUTING SERVICES</u>	)
<u>PROTOCOL</u>	)

**CERTIFICATE OF MAILING**

"Express Mail" mailing label no: EL575422828US

Date of Deposit: June 16, 2000

I hereby certify that this correspondence is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR

1.10 on the date indicated above and is addressed to:

Box Patent Application

Assistant Commissioner for Patents

Washington, D.C. 20231

Diane Morse

**TRANSMITTAL LETTER**

Honorable Assistant Commissioner  
for Patents  
Box Patent Application  
Washington, D.C. 20231

Dear Sir:

Enclosed for filing please find the patent application for an invention entitled, "CONTENT ROUTING SERVICES PROTOCOL", filed on behalf of Cisco Technology, Inc., assignee from inventor Douglas Gourlay, including Utility Patent Application Transmittal, 19 pages of specification, 17 pages of claims, 13 sheets of drawing figures, and 1 page of Abstract.


The attorney's Docket Number is CISCO-2369.

Kindly address all communications regarding this application to:

David B. Ritchie  
D'Alessandro & Ritchie  
P.O. Box 640640  
San Jose, CA 95164-0640  
Telephone (408) 441-1100

No fee is being paid at this time.

Respectfully submitted,  
D'ALESSANDRO & RITCHIE



---

Marc S. Hanish  
Reg. No. 42,626

Dated: June 16, 2000

D'Alessandro & Ritchie  
P.O. Box 640640  
San Jose, CA 95164-0640  
(408) 441-1100

**This application is submitted in the name of inventor Douglas Gourlay, assignor to Cisco Technology, Inc., a California Corporation.**

## S P E C I F I C A T I O N

### **CONTENT ROUTING SERVICES PROTOCOL**

#### BACKGROUND OF THE INVENTION

##### Field of the Invention

The present invention relates to the field of network communications. More specifically, the present invention relates to the use of a content routing service protocol to provide distributed Layer 7 switching in a computer network.

##### The Background

Most computer networked application architectures are organized as a series of layers, each layer built upon the one below it. The Open Systems Interconnect model is to have seven layers (L1-L7) in the system. Switching has historically taken place on Layers 2 and 3, with a recent emergence of Layer 4. In L2 switching, a local Media Access Control (MAC) address may be used to determine where to forward each packet. For this reason, L2 switching is not favored for large-scale networks such as the Internet or any LAN with more than two hundred and fifty nodes because of broadcast radiation issues.

L3 switching involves making a decision on where to forward each packet based on an IP (or similar) address located in the header of the packet. Due to its scalability,

L3 is currently the industry-standard for interconnecting networks and in the core of corporate LANs.

Layer 4 (L4) switching has been steadily gaining acceptance in the industry. L4 switching employs the information contained within the Transport Layer header to assist in switching and traffic handling. The L4 information indicates which protocol type is contained within each IP packet. Thus, L4 switching provides for several advantages: packet filtering, security, and quality of service. These are important features of L4 switching. However, the most crucial feature is server load balancing.

In server load balancing, a collection of physical servers, each with a different IP address, may support the same application service and be defined as a single virtual server. This virtual server becomes the single "logical server" with a single IP address. Therefore, rather than communicating directly with the real IP addresses of the physical servers, users direct traffic to the virtual server address.

In order to make this transition transparent to the user, a L4 switch may be used to direct the traffic to an appropriate server based on current traffic conditions. This also provides for the most effective use of the servers during high traffic conditions. FIG. 1 is a block diagram illustrating server load balancing. A server load balancer 2 in a L4 switch receives a syn packet from a user 4. The syn packet may have a destination address indicating the single virtual server. The server load balancer 2 then rewrites the

destination address in the syn packet as one of the servers' 6a, 6b, 6c IP address. In doing so, it may attempt to balance the current traffic load evenly among the servers 6a, 6b, 6c.

Layer 7 (L7), also known as application level switching, has been a goal of many companies for years. In L7 switching, the directing of packets may occur on a content level basis, where packets are directed to appropriate destinations based on the type of information they contain. Currently, some companies claim to have L7 switching capabilities, but their solutions involve routing packets based on destination IP address using a statically configured group of servers via a generic predictor algorithm. True L7 switching would allow content to be distributed throughout the Internet much in the same way some television programming is syndicated to individual stations rather than shown on a network.

FIG. 2 is a diagram illustrating a normal transaction involving a user's request for content. User 100 may wish to view content delivered by [www.contentprovider.com](http://www.contentprovider.com). In order to do this, the user 100 must directly contact [contentprovider.com](http://contentprovider.com)'s servers 102 to request the content. An IP syn packet must be sent to initiate a session, followed by an ack packet sent as a response [contentprovider.com](http://contentprovider.com)'s servers 102. Then a request for content must be sent, followed by [contentprovider.com](http://contentprovider.com)'s servers 102 returning the content to the user. Not only might this involve many "hops" on the Internet (and thus subject the request to traffic and other delays), but it also increases the danger of the user

not being able to view the content, such as if contentprovider.com's servers 102 are down, or a switch 104 connected to its servers 102 is down.

If there was a solution that provided for true content-based routing, the content  
5 could be distributed at different points on the Internet and a direct connection to contentprovider.com's servers 102 need not occur. With the explosive growth of the Internet, it would be beneficial to have such a solution.

SUMMARY OF THE INVENTION

Layer 7 switching may be accomplished using one or more caches placed throughout a computer network. Changes to a file on a server may be detected and propagated throughout the network. At the switch or router level, once notification of changes to a file is received, the content may be retrieved from the server and placed in a connected cache. A routing table entry may be created for the content and also placed in the cache. The routing table entry may contain an original location field identifying the original location of the content, a distance field indicating a distance from the cache to the server, and a field indicating a version number of the content. Additional fields may also be contained within the routing table entry. When a user requests a specific file, rather than forward the request directly to the server containing the original file, the request may be handled by the router closest to the user which has a connected cache containing the content. This allows a user's request to be handled much more quickly and efficiently than prior art solutions.



## BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram illustrating server load balancing.

FIG. 2 is a diagram illustrating a normal transaction involving a user's request for  
5 content.

FIG. 3 is a block diagram illustrating a system for content routing in a computer network in accordance with a specific embodiment of the present invention.

FIG. 4 is a diagram illustrating the format of each record in accordance with a  
10 specific embodiment of the present invention.

FIG. 5 is a diagram illustrating how content changes may be propagated through  
the system for content switching in accordance with a specific embodiment of the present  
15 invention.

FIG. 6 is a diagram illustrating an example of a routing table entry.

FIG. 7 is a diagram illustrating how the system for content switching in  
20 accordance with a specific embodiment of the present invention would appear after  
content changes have been propagated throughout the network.

FIG. 8 is a flow diagram illustrating a method for making content available for users in a computer network in accordance with a specific embodiment of the present invention.

5 FIG. 9 is a flow diagram illustrating a method for updating content in a computer network, the content located at a web server and having an original location, in accordance with another embodiment of the present invention.

10 FIG. 10 is a flow diagram illustrating a method for handling a request for content from a user in a computer network in accordance with another embodiment of the present invention.

15 FIG. 11 is a block diagram illustrating an apparatus for making content available for users in a computer network in accordance with a specific embodiment of the present invention.

FIG. 12 is a block diagram illustrating an apparatus for updating content in a computer network, the content located at a web server and having an original location, in accordance with another embodiment of the present invention.

## DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

In the following description, a preferred embodiment of the invention is described with regard to preferred process steps and data structures. However, those skilled in the art will recognize, after perusal of this application, that embodiments of the invention may be implemented using at least one general purpose computer operating under program control, and that modification of the general purpose computer to implement the components, process steps, and/or data structures described herein would not require undue invention.

In accordance with a presently preferred embodiment of the present invention, the components, process steps, and/or data structures are implemented using software. This implementation is not intended to be limiting in any way. Different implementations may be used and may include other types of operating systems, computing platforms, program storage devices and/or computer programs. In addition, those of ordinary skill in the art will readily recognize that devices of a less general purpose nature, such as hardwired devices, devices relying on FPGA (field programmable gate array) or ASIC (application specific integrated circuit) technology, or the like, may also be used without departing from the scope and spirit of the inventive concepts disclosed herewith.

FIG. 3 is a block diagram illustrating a system for content routing in a computer network in accordance with a specific embodiment of the present invention. Switches

150a - 150i permeate the system. Attached to one or more of the switches 150a-150i is a cache 152a - 152f. The caches 152a - 152f also may store the content locally to improve speed and bandwidth. The caches 152a - 152f may each store a routing table. The routing table may have a record for each piece of content stored in the cache. The content may then be identified by an http url address. However, the address may point to a specific file (such as "www.contentprovider.com/logos/logo.gif") or to a higher-level directory (such as "www.contentprovider.com/logos", indicating all files and subdirectories contained within the higher level directory).

FIG. 4 is a diagram illustrating the format of each record in the routing table accordance with a specific embodiment of the present invention. A URL field 200 contains an http-style address as discussed above. An IP/Port field 202 may indicate the IP address and/or port from which the content originated. A Metric field 204 may indicate the distance from the switch to the content. Bandwidth and delay could be initial metrics. A Serial Number field 206 may indicate a serial number or version number for the content. A Billing token field 208 may be used to exchange tokens between service providers or between a service provider and a content provider. This allows the use of content to be tracked by a single identifier per content provider. This also allows for an 800 number style billing system.

A Date/Time stamp field 210 records the date and time of the last update to the record, which is important in synchronizing records across routers. Finally, an additional

tag field 212 is provided for other functions, such as how to handle an object. Various options have been contemplated, such as storing a route in a routing table permanently for a more static-type routing, updating content serving sites, bypassing the cache, or applying a Quality of Service (QoS) parameter based on billing tags, urls, etc.

5

Referring back to FIG. 3, when a user wishes to view content created by contentprovider.com, as usual he will generate an http request packet with the corresponding http url (for example, "www.contentprovider.com/logos/logo.gif"). A normal http request packet requires a three-way handshaking procedure with a switch or router. SYN and ACK fields in the headers of the packets sent between the switch or router and the user indicate the stages of handshaking. When SYN=1 and ACK=0, the packet is an "open connection" packet sent by the user. When SYN =1 and ACK=1 the packet is an "open connection acknowledgement" packet sent by the switch or router. Finally, when SYN =0 and ACK = 1 the packet indicates that the user is ready to send data. This results in a delay until at least the 4th packet is processed before the switch or router can examine information in the header of the payload packets.

10  
15

Eliminating the delay would allow a switch or router to determine whether or not it wishes to even establish a TCP/IP session based on content-specific information. This packet is forwarded to the closest switch 150a. However, rather than simply forward the request packet to another switch with the intention of eventually getting the packet to contentprovider.com's servers 154 for processing, the switch 150a may examine the

20

routing table stored in its cache, assuming it has a cache, searching for an entry with the identical url. If one with an identical url exists, then the content resides in the local cache and the request packet need not be forwarded to contentprovider.com's servers 154.

Additionally, a TCP/IP session may then be established to fetch the content from the

5 local cache and send it to the user. Should the nearest switch's cache not have an entry for that url (or the switch not contain a cache at all), the request packet may simply be forwarded to the next closest switch 150b en route to contentprovider.com's servers without establishing a TCP/IP session between the previous switch or router and the user.

Then the next switch may check its local cache much in the same way. Thus, the worst-

10 case scenario is that none of the switches en route to contentprovider.com's servers contain the content in their local caches, and the content need only be retrieved directly at contentprovider.com's servers 154. However, if content distribution is done properly,

then many times this worst-case scenario can be avoided and the content may be retrieved from a cache located at a point much closer to the user than contentprovider.com's servers

15 154. The distance from a particular cache to the content provider's servers may be determined by examining the distance field in the record for the content in the cache's routing table.

Thus, the speed of accessing any particular piece of content has been greatly

20 enhanced. Additionally, the reliability has also been enhanced, as users will be able to access contentprovider.com content even if contentprovider.com's servers 154 are down or unreachable.

The http request packet may take a different form than it has in the past in order to facilitate this design. Rather than send a SYN=1, ACK=0 packet (also called a syn packet) alone to initiate the session, a special payload called a HUP packet will be

5 appended to the syn packet. The HUP packet will contain the full url of the http address being requested. This packet will take the form A-C-B, where A is the top-level domain name (e.g. "www.contentprovider.com"), C is the file name (e.g. "logo.gif"), and B is the rest of the address (e.g. "/logos/"). In this way, the router or switch has the ability to determine how to handle the request based on just the top level domain name (e.g. "this

10 router does not have any content from contentprovider.com, thus forward the request elsewhere) or based on just the top level domain name and the file type (e.g. "this router has been instructed to forward all requests for .gif files to a special cache").

Turning now to the distribution of the content throughout the network, FIG. 5 is a

15 diagram illustrating how content changes may be propagated through the system for content switching in accordance with a specific embodiment of the present invention. A content based web site 250, such as contentprovider.com, may provide various graphics, web pages, and other files as content to be distributed to various users throughout the Internet. For purposes of this example, the url [www.contentprovider.com/logos/logo.gif](http://www.contentprovider.com/logos/logo.gif)

20 is used to represent one of the content files contentprovider.com provides. Version 2 of this logo may be already stored in caches throughout the network.

FIG. 6 is a diagram illustrating an example of a routing table entry. As can be seen, the entry for "http://www.contentprovider.com/logos/logo.gif" 300 has serial/version number 2 302. When the logo is changed on contentprovider.com to version 3, there is a situation where the content needs to be updated in multiple sites throughout the network. Referring back to FIG. 5, the server 250 may generate a route to the nearest switch (alternatively, a server load balancing platform may detect the content change and advertise the change to an upstream peer). The nearest switch may then examine its local routing table, looking for an entry sharing the same original location url in the url field. If one exists, the version numbers may be compared to determine if the content needs to be update. Assuming the information received indicates a higher version number than the one already stored in the cache, the cache will undertake to update the content, including requesting the new version of the content from the server 250, storing the new content in the cache, updating its local routing table with the new version number, and updating the time/date stamp. The route may be propagated throughout the network until the logo with version 3 is known on all switches/routers.

At any particular cache, once the content update is received, the routing table entry may be updated. This will generally include at least changing the date/time stamp and the serial/version number.

FIG. 7 is a diagram illustrating how the system for content switching in accordance with a specific embodiment of the present invention would appear after



content changes have been propagated throughout the network. Version 3 of the content is now stored in one or more of the caches 400a-400e distributed throughout the network. Additionally, routing table entries for the content at the one or more caches 400a-400e now contain the updated version number and time stamp.

5

Other updates to the routing tables may be accomplished via a similar process. Examples of other updates include deletion of the content, addition of new content, alteration of the tag field, and alteration of the billing token field.

10  
15  
20

FIG. 8 is a flow diagram illustrating a method for making content available for users in a computer network in accordance with a specific embodiment of the present invention. At 450 the content is forwarded to one or more caches distributed throughout the computer network, each of the caches coupled to a switch or router. At 452 the content is stored in each of the one or more caches. Finally, at 454, a record identifying the content is stored in each of said one or more caches, said record for any particular cache of said one or more caches having an original location field identifying the original location of said content, a distance field indicating a distance from said particular cache to said original location of said content, and a version number field indicating a version number of said content.

20

FIG. 9 is a flow diagram illustrating a method for updating content in a computer network, the content located at a web server and having an original location, in

accordance with another embodiment of the present invention. The original location of the content may be in the form of an http url. At 500, the content is altered. For example, a logo may be changed to a newer version. At 502, the alteration of the content is detected with a server load balancer through polling of the web server. Alternatively,

5 the web server may directly notify the server load balancer when an updated has been performed. At 504, a routing table entry is created for the content in a cache, said routing table entry having a record with a location field with the original location of said content, a distance field indicating the distance from said cache to the original location of said content, and a version number field indicating a version number of said content. The

10 record may also include an IP address field that indicates the IP address and/or port of the web server. Fields indicating the date and time of the last update to the record, billing tokens or certificates for content peering between providers, Quality of Service processes to be applied when a user attempts to access the content, whether the content need not be stored in the cache, and server load balancing processes to be applied when a user

15 attempts to access the content may also be added.

At 506, said routing table entry is forwarded to another of one or more caches in the computer network. The creating and forwarding are then repeated for each of said one or more caches. The distance from a to the original location may be recomputed

20 when the routing table entry is forwarded to the cache. Additionally, for each cache it may be determined whether a record corresponding to an older version of the content is already stored in the cache and the content may be retrieved from the original location

and stored in the cache if a record corresponding to an older version of the content is already stored in the cache.

FIG. 10 is a flow diagram illustrating a method for handling a request for content from a user in a computer network in accordance with another embodiment of the present invention. At 550, the request is received at a switch or router. At 552, an original location address in a header in the request is examined. At 554, it is determined if the cache is coupled to the switch or router. At 556, the original location address is compared with one or more entries in a table in a cache coupled to the switch or router, if such a cache exists. At 558, it is determined if an entry in the table in the cache has an original location field identical to the original location address. At 560, the content is forwarded from the cache to the user if an entry in the table in the cache has an original location field identical to said original location address. On the other hand, if the cache does not exist or the cache does not have an entry in the table with an original location field identical to the original location address, at 562 the request is transferred to another switch or router. The switch or router it is forwarded to should be one that further along in a path ending with a web server hosting the content. This ensures that if no caches contain the content it will still be possible for the user to retrieve the content directly from the web server.

FIG. 11 is a block diagram illustrating an apparatus for making content available for users in a computer network in accordance with a specific embodiment of the present

invention. A content forwarder 600 coupled to a server load balancer 602 forwards content through the computer network 604 to one or more caches 606 distributed throughout the computer network, each of the caches coupled to a switch or router 608. A content storer 610 coupled to said cache 606 and said switch or router 608 stores content in the cache 606. A record storer 612 coupled to said cache 606 and said switch or router 608 stores a record in said cache 606, the record identifying the content in each of said one or more caches, said record for any particular cache of said one or more caches having an original location field identifying the original location of said content, a distance field indicating a distance from said particular cache to said original location of said content, and a version number field indicating a version number of said content.

FIG. 12 is a block diagram illustrating an apparatus for updating content in a computer network, the content located at a web server and having an original location, in accordance with another embodiment of the present invention. The original location of the content may be in the form of an http url. The content is altered at a server. For example, a logo may be changed to a newer version. The alteration of the content is detected with a server load balancer through polling of the web server. Alternatively, the web server may directly notify the server load balancer when an update has been performed. A routing table entry creator 650 creates a routing table entry for the content in a cache 652, said routing table entry having a record with a location field with the original location of said content, a distance field indicating the distance from said cache 652 to the original location of said content, and a version number field indicating a

version number of said content. The record may also include an IP address field that indicates the IP address and/or port of the web server. Fields indicating the date and time of the last update to the record, billing tokens or certificates for content peering between providers, Quality of Service processes to be applied when a user attempts to access the content, whether the content need not be stored in the cache, and server load balancing processes to be applied when a user attempts to access the content may also be added.

A routing table entry forwarder 654 coupled to said routing table entry creator forwards the routing table entry to another of one or more caches in the computer network. The creating and forwarding are then repeated for each of said one or more caches. The distance from a cache 652 to the original location may be recomputed when the routing table entry is forwarded to the cache using a cache-to-original-location distance recomputer 656. Additionally, for each cache it may be determined using a record version determiner 658 whether a record corresponding to an older version of the content is already stored in the cache and the content may be retrieved from the original location and stored in the cache using a content retriever 660 if a record corresponding to an older version of the content is already stored in the cache.

FIG. 13 is a block diagram illustrating an apparatus for handling a request for content from a user in a computer network in accordance with another embodiment of the present invention. A request receiver 700 receives the request at a switch or router. An original location address examiner 702 coupled to said request receiver 700 and to a

cache 704 examines an original location address in a header in the request. An original location address comparator 706 coupled to said original location address examiner 702 and to said cache 704 compares the original location address with one or more entries in a table in said cache 704. A content forward 708 coupled to said original location address  
 5 comparator and to said cache 704 forwards the content from the cache 704 to the user if an entry in the table in the cache 704 has an original location field identical to said original location address. On the other hand, if the cache does not exist or the cache does not have an entry in the table with an original location field identical to the original location address, the request is transferred to another switch or router using a request  
 10 transferer 710. The switch or router it is forwarded to should be one that further along in a path ending with a web server hosting the content. This ensures that if no caches contain the content that it will still be possible for the user to retrieve the content directly from the web server.

15 While embodiments and applications of this invention have been shown and described, it would be apparent to those skilled in the art having the benefit of this disclosure that many more modifications than mentioned above are possible without departing from the inventive concepts herein. The invention, therefore, is not to be restricted except in the spirit of the appended claims.

CLAIMS

What is claimed is:

- 5 1. A method for making content available for users in a computer network, the content having an original location, the method including:
  - forwarding the content to one or more caches distributed throughout the computer network, each of said caches coupled to a switch or router;
  - storing the content in each of said one or more caches; and
  - 10 storing a record identifying said content in each of said one or more caches, said record for any particular cache of said one or more caches having an original location field identifying the original location of said content, a distance field indicating a distance from said particular cache to the original location of said content, and a field indicating a version number of said content.
- 15 2. The method of claim 1, wherein said original location of said content includes a hypertext transfer protocol uniform resource locator.
3. The method of claim 1, wherein said record further includes a field indicating an IP address for a web server hosting said content.
- 20

4. The method of claim 3, wherein said field indicating an IP address further indicates a port number.

5. The method of claim 1, wherein said record further includes a field indicating a  
5 date and time of the last update to the record.

6. The method of claim 1, wherein said record further includes a field having a billing token or certificate for content peering between providers.

10 7. The method of claim 1, wherein said record further includes a tag field indicating a Quality of Service process to be applied when a user attempts to access the content.

8. The method of claim 1, wherein said record further includes a cache bypass field indicating whether said content need not be stored in said particular cache.

15 9. The method of claim 1, wherein said record further includes a server load balancing field indicating a server load balancing process to be applied when a user attempts to access the content.

20 10. A method for making content available for users in a computer network, the content having an original location, the method including:

receiving the content forwarded from the original location;



storing the content in a cache coupled to a switch or router; and

storing a record identifying the content in said cache, said record having an original location field identifying the original location of said content, a distance field indicating a distance from said cache to the original location of said content, and a field  
5 indicating a version number of said content.

11. The method of claim 10, wherein said original location of said content includes a hypertext transfer protocol uniform resource locator.

10 12. The method of claim 10, wherein said record further includes a field indicating an IP address for a web server hosting said content.

13. The method of claim 12, wherein said field indicating an IP address further indicates a port number.  
15

14. The method of claim 10, wherein said record further includes a field indicating a date and time of the last update to the record.

15. The method of claim 10, wherein said record further includes a field having a  
20 billing token or certificate for content peering between providers.

16. The method of claim 10, wherein said record further includes a tag field indicating a Quality of Service process to be applied when a user attempts to access the content.

17. The method of claim 10, wherein said record further includes a cache bypass field  
5 indicating whether said content need not be stored in said cache.

18. The method of claim 10, wherein said record further includes a server load balancing field indicating a server load balancing process to be applied when a user attempts to access the content.

19. A method for updating content in a computer network, the content located at a web server and having an original location, the method including:

creating a routing table entry for the content in a cache, said routing table entry having a record with a location field with the original location of said content, a distance  
5 field indicating the distance from said cache to the original location of said content, and a version number field indicating a version number of said content; and

forwarding said routing table entry to another of one or more caches in the computer network to allow said another of one or more caches to create a routing table entry for the content.

20. The method of claim 19, wherein said field with the original location of the content includes a hypertext transfer protocol uniform resource locator.

21. The method of claim 19, wherein said record further includes a field indicating an IP address for the web server.

5 22. The method of claim 21, wherein said field indicating the IP address further indicates a port number.

23. The method of claim 19, wherein said record further includes a field indicating the date and time of the last update to the record.

24. The method of claim 19, wherein said record further includes a field having a billing token or certificate for content peering between providers.

25. The method of claim 19, wherein said record further includes a tag field indicating a Quality of Service process to be applied when a user attempts to access the content.

26. The method of claim 19, wherein said record further includes a cache bypass field indicating whether said content need not be stored in said cache.

20 27. The method of claim 19, wherein said record further includes a server load balancing field indicating a server load balancing process to be applied when a user attempts to access the content.

28. The method of claim 19, further including:

determining if a record corresponding to an older version of said content is already stored in said cache; and

5 retrieving the content from said original location and storing it in said cache if a record corresponding to an older version of said content is already stored in said cache.

29. The method of claim 20, wherein said determining include examining said version number field of said record and comparing it with version number fields in records with identical location fields.

30. The method of claim 19, further including:

detecting the alteration of the content with a server load balancer through polling of the web server.

31. The method of claim 19, further including:

notifying a server load balancer that the content has been altered.

32. A method for handling a request for content from a user in a computer network,

20 including:

receiving the request at a switch or router;

examining an original location address in a header in the request;

comparing said original location address with one or more entries in a table in a cache coupled to said switch or router, if said cache exists;

forwarding the content from said cache to said user if an entry in said table in said cache has an original location field identical to said original location address; and

5 transferring said request to another switch or router if said cache does not exist or said cache does not have an entry in said table with an original location field identical to said original location address.

33. The method of claim 32, wherein said transferring includes transferring said request to the next switch or router along a path ending with a web server hosting the content.

34. The method of claim 32, wherein said request is in the form of a SYN packet with said original address located in a header in a payload field of the SYN packet.

35. The method of claim 34, establishes a TCP/IP session between itself and the user if an entry in said table in said cache has an original location field identical to said original location address.

20 36. An apparatus for making content available for users in a computer network, the content having an original location, the apparatus including:

a content forwarder coupled to a server load balancer;

a switch or router coupled to said content forwarder via the computer network;  
 a cache coupled to said switch or router;  
 a content storer coupled to said cache and coupled to said switch or router; and  
 a record storer coupled to said cache and coupled to said switch or router.

5

37. An apparatus for updating content in a computer network, the content located at a web server and having an original location, the apparatus including:

a cache;  
 a routing table entry creator coupled to said cache; and  
 a routing table entry forwarder coupled to said cache and to said routing table entry creator.

38. The apparatus of claim 37, further including:

a record version determiner coupled to said cache; and  
 a content retriever coupled to said record version determiner.

39. The apparatus of claim 37, further including a cache-to-original-location distance recomputer coupled to said routing table entry creator.

40. An apparatus for handling a request for content from a user in a computer network, including:

a request receiver;

a cache;

an original location address examiner coupled to said request receiver and to said cache;

an original location address comparator coupled to said original location address  
5 examiner and to said cache receiving the request at a switch or router;

a content forwarder coupled to said original location address comparator and to said cache; and

a request transferer coupled to said request receiver and to said original location address comparator.

41. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine for making content available for users in a computer network, the content having an original location, the method including:

forwarding the content to one or more caches distributed throughout the computer  
15 network, each of said caches coupled to a switch or router;

storing the content in each of said one or more caches; and

storing a record identifying said content in each of said one or more caches, said record for any particular cache of said one or more caches having an original location field identifying the original location of said content, a distance field indicating a distance  
20 from said particular cache to the original location of said content, and a field indicating a version number of said content.

42. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine for making content available for users in a computer network, the content having an original location, the method including:

receiving the content forwarded from the original location;

5 storing the content in a cache coupled to a switch or router; and

storing a record identifying the content in said cache, said record having an original location field identifying the original location of said content, a distance field indicating a distance from said cache to the original location of said content, and a field indicating a version number of said content.

43. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine for updating content in a computer network, the content located at a web server and having an original location, the method including:

creating a routing table entry for the content in a cache, said routing table entry

5 having a record with a location field with the original location of said content, a distance field indicating the distance from said cache to the original location of said content, and a version number field indicating a version number of said content;

forwarding said routing table entry to another of one or more caches in the computer network; and

20 repeating said creating and forwarding for each of said one or more caches.



44. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine for updating content in a computer network, the content located at a web server and having an original location, the method including:

creating a routing table entry for the content in a cache, said routing table entry

5 having a record with a location field with the original location of said content, a distance field indicating the distance from said cache to the original location of said content, and a version number field indicating a version number of said content; and

forwarding said routing table entry to another of one or more caches in the computer network to allow said another of one or more caches to create a routing table entry for the content.

45. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine for handling a request for content from a user in a computer network, the method including:

15 receiving the request at a switch or router;

examining an original location address in a header in the request;

comparing said original location address with one or more entries in a table in a cache coupled to said switch or router, if said cache exists;

forwarding the content from said cache to said user if an entry in said table in said

20 cache has an original location field identical to said original location address; and

transferring said request to another switch or router if said cache does not exist or said cache does not have an entry in said table with an original location field identical to said original location address.

5

46. An apparatus for making content available for users in a computer network, the content having an original location, the apparatus including:

a memory;

a network interface coupled to said memory;

a processor coupled to said memory programmed to perform the steps of:

forwarding the content to one or more caches distributed throughout the computer network, each of said caches coupled to a switch or router;

storing the content in each of said one or more caches; and

storing a record identifying said content in each of said one or more caches, said record for any particular cache of said one or more caches having an original location field identifying the original location of said content, a distance field indicating a distance from said particular cache to the original location of said content, and a field indicating a version number of said content.

20 47. An apparatus for making content available for users in a computer network, the content having an original location, the apparatus including:

a memory;

a network interface coupled to said memory;

a processor coupled to said memory programmed to perform the steps of:

receiving the content forwarded from the original location;

storing the content in a cache coupled to a switch or router; and

5 storing a record identifying the content in said cache, said record having an original location field identifying the original location of said content, a distance field indicating a distance from said cache to the original location of said content, and a field indicating a version number of said content.

0 48. An apparatus for updating content in a computer network, the content located at a web server and having an original location, the apparatus including:

a memory;

a network interface coupled to said memory;

a processor coupled to said memory programmed to perform the steps of:

15 creating a routing table entry for the content in a cache, said routing table entry having a record with a location field with the original location of said content, a distance field indicating the distance from said cache to the original location of said content, and a version number field indicating a version number of said content;

20 forwarding said routing table entry to another of one or more caches in the computer network; and

repeating said creating and forwarding for each of said one or more caches.

49. An apparatus for updating content in a computer network, the content located at a web server and having an original location, the apparatus including:

a memory;

5 a network interface coupled to said memory;

a processor coupled to said memory programmed to perform the steps of:

creating a routing table entry for the content in a cache, said routing table entry having a record with a location field with the original location of said content, a distance field indicating the distance from said cache to the original location of said content, and a version number field indicating a version number of said content; and

forwarding said routing table entry to another of one or more caches in the computer network to allow said another of one or more caches to create a routing table entry for the content.

50. An apparatus for handling a request for content from a user in a computer network, including:

a memory;

a network interface coupled to said memory;

20 a processor coupled to said memory programmed to perform the steps of:

receiving the request at a switch or router;

examining an original location address in a header in the request;

comparing said original location address with one or more entries in a table  
in a cache coupled to said switch or router, if said cache exists;

forwarding the content from said cache to said user if an entry in said table  
in said cache has an original location field identical to said original location  
5 address; and

transferring said request to another switch or router if said cache does not  
exist or said cache does not have an entry in said table with an original location  
field identical to said original location address.

51. An apparatus for making content available for users in a computer network, the  
content having an original location, the apparatus including:

means for forwarding the content to one or more caches distributed throughout the  
computer network, each of said caches coupled to a switch or router;

means for storing the content in each of said one or more caches; and

5 means for storing a record identifying said content in each of said one or more  
caches, said record for any particular cache of said one or more caches having an original  
location field identifying the original location of said content, a distance field indicating a  
distance from said particular cache to the original location of said content, and a field  
indicating a version number of said content.

20 52. An apparatus for making content available for users in a computer network, the  
content having an original location, the apparatus including:

means for receiving the content forwarded from the original location;

means for storing the content in a cache coupled to a switch or router; and

means for storing a record identifying the content in said cache, said record having an original location field identifying the original location of said content, a distance field  
 5 indicating a distance from said cache to the original location of said content, and a field indicating a version number of said content.

53. An apparatus for updating content in a computer network, the content located at a web server and having an original location, the apparatus including:

10 means for creating a routing table entry for the content in a cache, said routing table entry having a record with a location field with the original location of said content, a distance field indicating the distance from said cache to the original location of said content, and a version number field indicating a version number of said content;

15 means for forwarding said routing table entry to another of one or more caches in the computer network; and

means for repeating said creating and forwarding for each of said one or more caches.

54. An apparatus for updating content in a computer network, the content located at a  
 20 web server and having an original location, the apparatus including:

means for creating a routing table entry for the content in a cache, said routing table entry having a record with a location field with the original location of said content,

a distance field indicating the distance from said cache to the original location of said content, and a version number field indicating a version number of said content; and

means for forwarding said routing table entry to another of one or more caches in the computer network to allow said another of one or more caches to create a routing  
5 table entry for the content.

55. An apparatus for handling a request for content from a user in a computer network, including:

means for receiving the request at a switch or router;

means for examining an original location address in a header in the request;

means for comparing said original location address with one or more entries in a table in a cache coupled to said switch or router, if said cache exists;

means for forwarding the content from said cache to said user if an entry in said table in said cache has an original location field identical to said original location address;  
15 and

means for transferring said request to another switch or router if said cache does not exist or said cache does not have an entry in said table with an original location field identical to said original location address.

ABSTRACT

Layer 7 switching may be accomplished using one or more caches placed throughout a computer network. Changes to a file on a server may be detected and propagated throughout the network. At the switch or router level, once notification of changes to a file is received, the content may be retrieved from the server and placed in a connected cache. A routing table entry may be created for the content and also placed in the cache. The routing table entry may contain an original location field identifying the original location of the content, a distance field indicating a distance from the cache to the server, and a field indicating a version number of the content. Additional fields may also be contained within the routing table entry. When a user requests a specific file, rather than forward the request directly to the server containing the original file, the request may be handled by the router closest to the user which has a connected cache containing the content. This allows a user's request to be handled much more quickly and efficiently than prior art solutions.



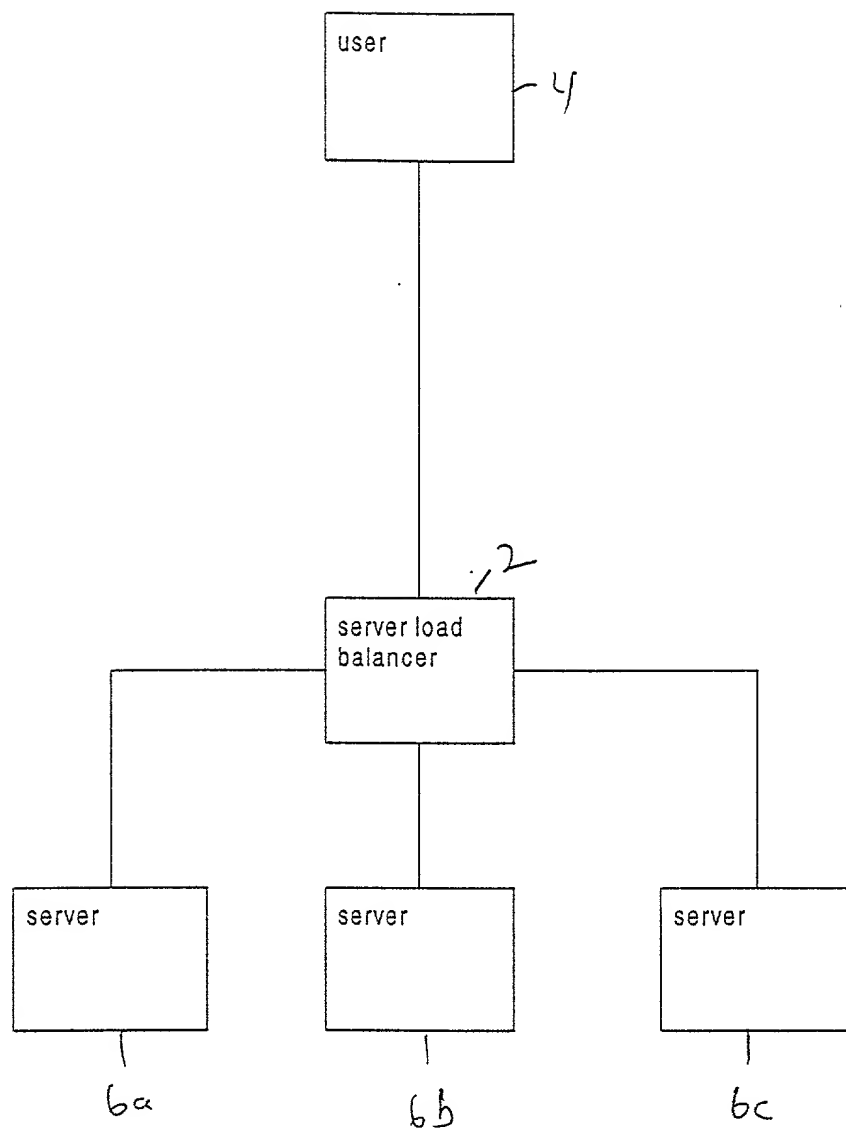


FIG. 1

Copyright © 2008 Cisco and/or its affiliates. All rights reserved. Cisco Confidential




FIG. 2

FIG. 16.3 is a schematic diagram of a network architecture for TCP termination. The diagram shows a client device 150a connected to a network of servers and routers. The network includes a central router 150b, which is connected to several other routers (150c, 150d, 150e, 150f, 150g, 150h, 150i) and servers (152a, 152b, 152c, 152e, 152f). The servers are connected to a content provider 154. The network is divided into three main sections: a left section with routers 150a, 150b, and 150c; a middle section with routers 150d, 150e, and 150f; and a right section with routers 150g, 150h, and 150i. The content provider 154 is located in the right section. The diagram illustrates the flow of data from the client device 150a through the network to the content provider 154, with the label 'TCP Termination' indicating the process of terminating TCP connections at the network level.




FIG. 16.3

URL	- 200
IP/Port	- 202
Metric	- 204
Serial/Version #	- 206
Billing Token	- 208
Date/Time Stamp	- 210
Tag	- 212



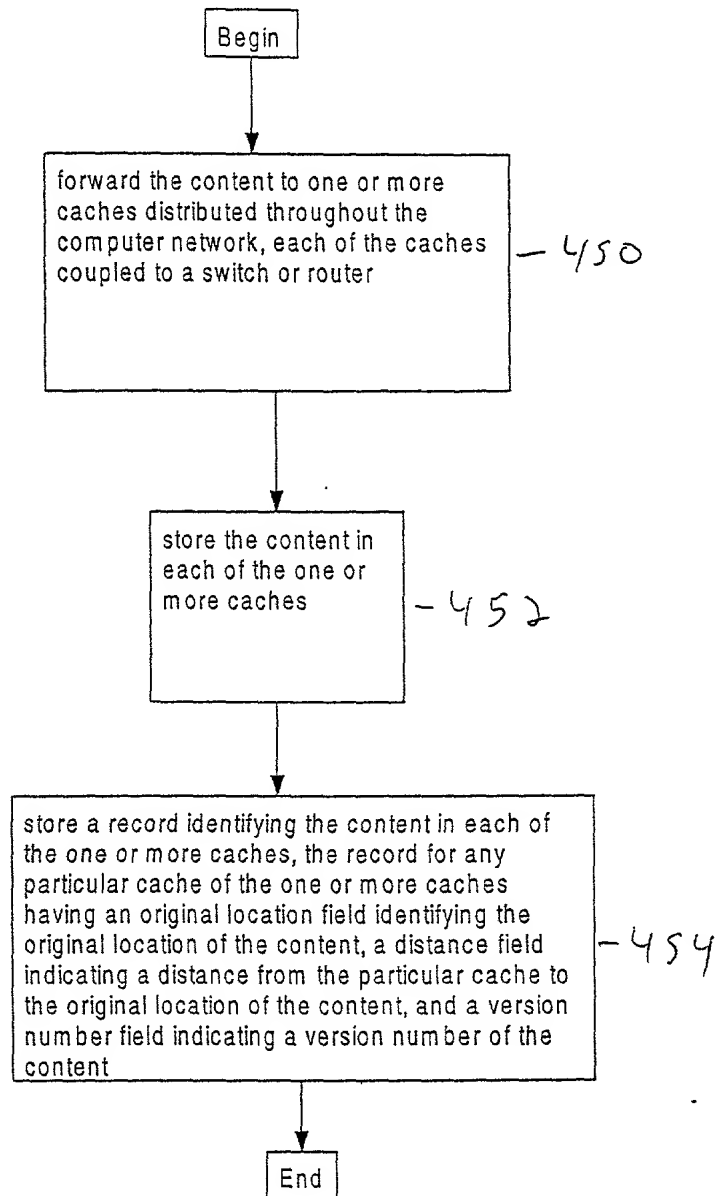
URL	http://www.cnn.com/logos/logo.gif
IP/Port	171.68.62.5:80
Metric	2
Serial/Version #	2
Billing Token	6687945
Date/Time Stamp	3-17-2000:23:04:56
Tag	16

- 300

- 302

F16.6







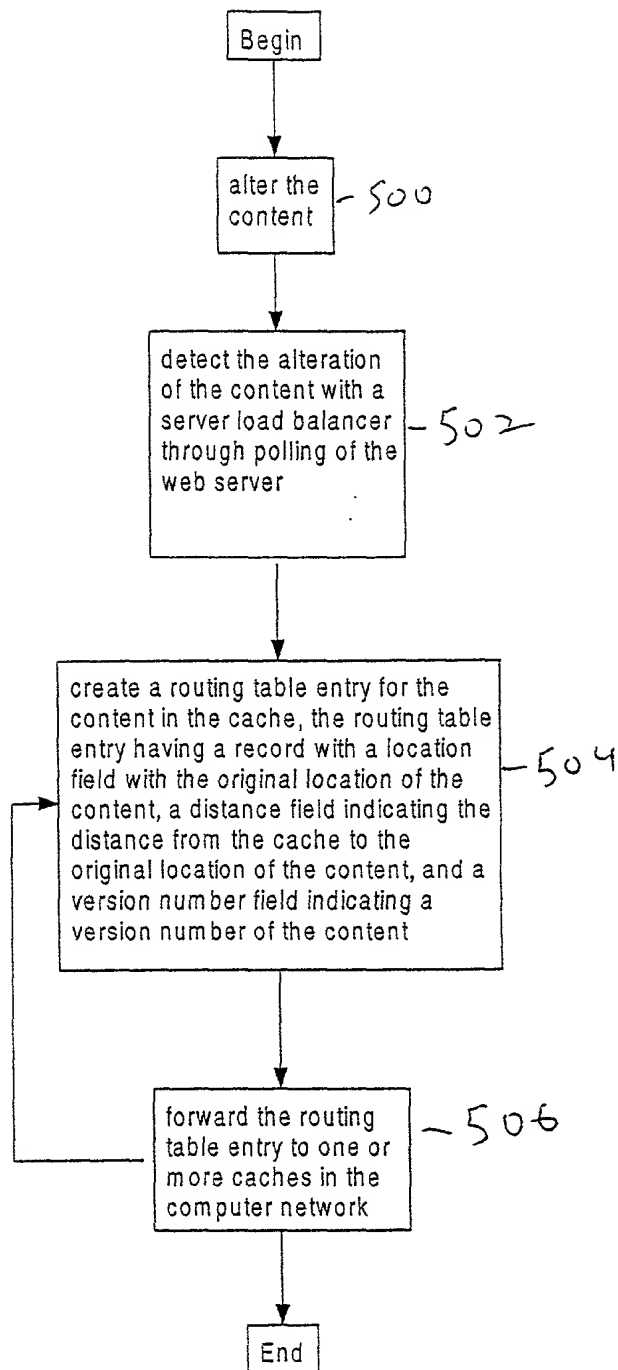


FIG. 9

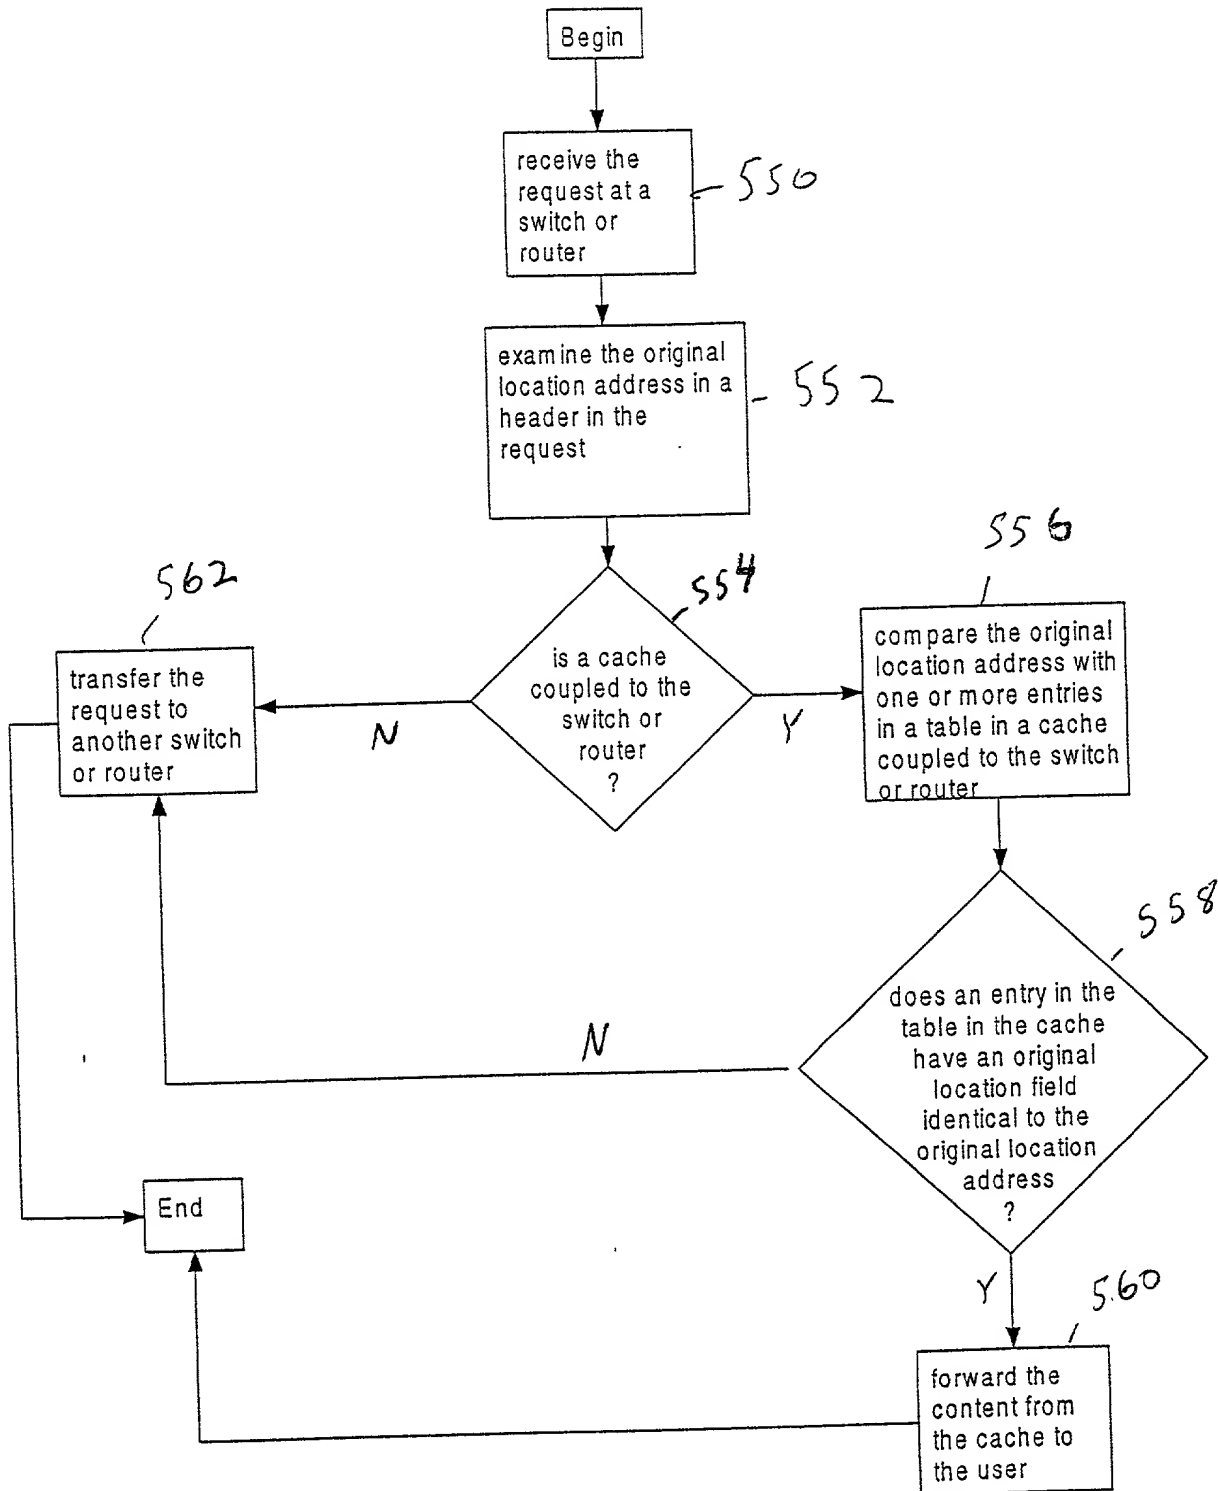


FIG. 10




FIG. 11

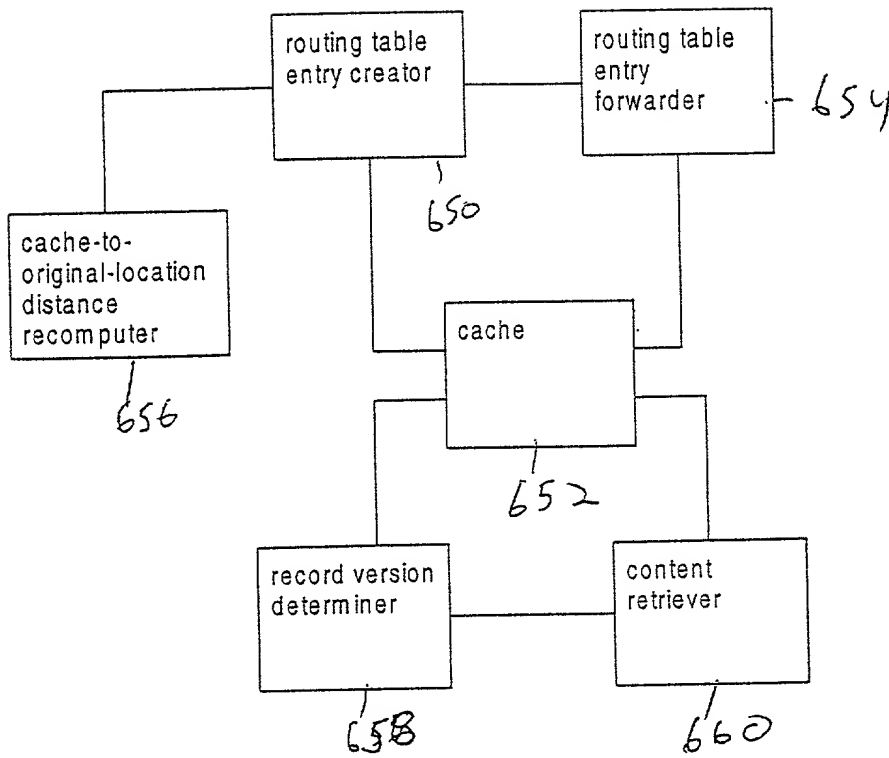


FIG. 12

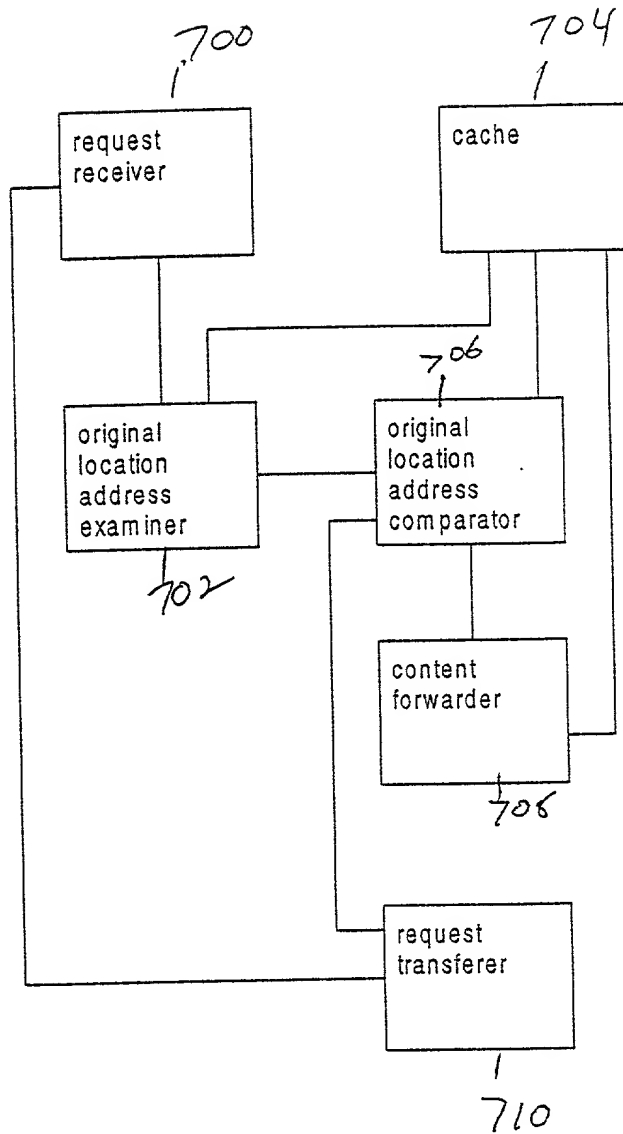


FIG. 13